The HappyFace Project

Site Specific Monitoring of Multiple Information Systems

Volker Büge
for the HappyFace Developers

Institut für Experimentelle Kernphysik
Karlsruhe Institute of Technology

Outline

- Motivation for a Meta-Monitoring System
  - Current monitoring situation
  - Advantages of meta-monitoring
- The HappyFace Project
  - Motivation
  - Architecture
  - Instances
  - Examples
- Summary & Outlook

Monitoring Situation

Getting all required information for a grid site is complicated

- Monitoring information is not clearly arranged, there are…
  - Many sources of valuable information
  - Different information displays provided by different technologies
- Totality of all monitoring systems is uncomfortable to use, you have to…
  - Manage many browser tabs / windows
  - Change the settings of the web interfaces (time range, site, …)
  - Long waiting time until page opens up, often more than 30 seconds
- Consequences:
  - Unnecessary increase of administration effort for a grid site
  - Difficult to identify correlations
  - Nearly impossible to get quick overview on a site's status for non experts, especially if several services at different sites are involved

Solution: Access to all needed information via one web page

Properties of a Meta Monitoring System

Idea to ease administration: Build up a meta monitoring system

- Such a framework should …
  - Collect and process all important monitoring information
  - Present the current status of a grid site and its services
  - Display simple rating / warning system (smiley faces, arrows, …)

- Design properties:
  - Framework has a modular layout: There is a static core that provides the basic functionality for the dedicated tests. The individual tests can be plugged in.
  - Decoupling of collecting the information and the actual visualisation
  - All information is accessible via a single website, including a history
  - Visualisation should provide a smart and quick overview on the monitored service which also allows to identify correlations

The HappyFace Project provides such a smart summary of existing information

The HappyFace Project Architecture

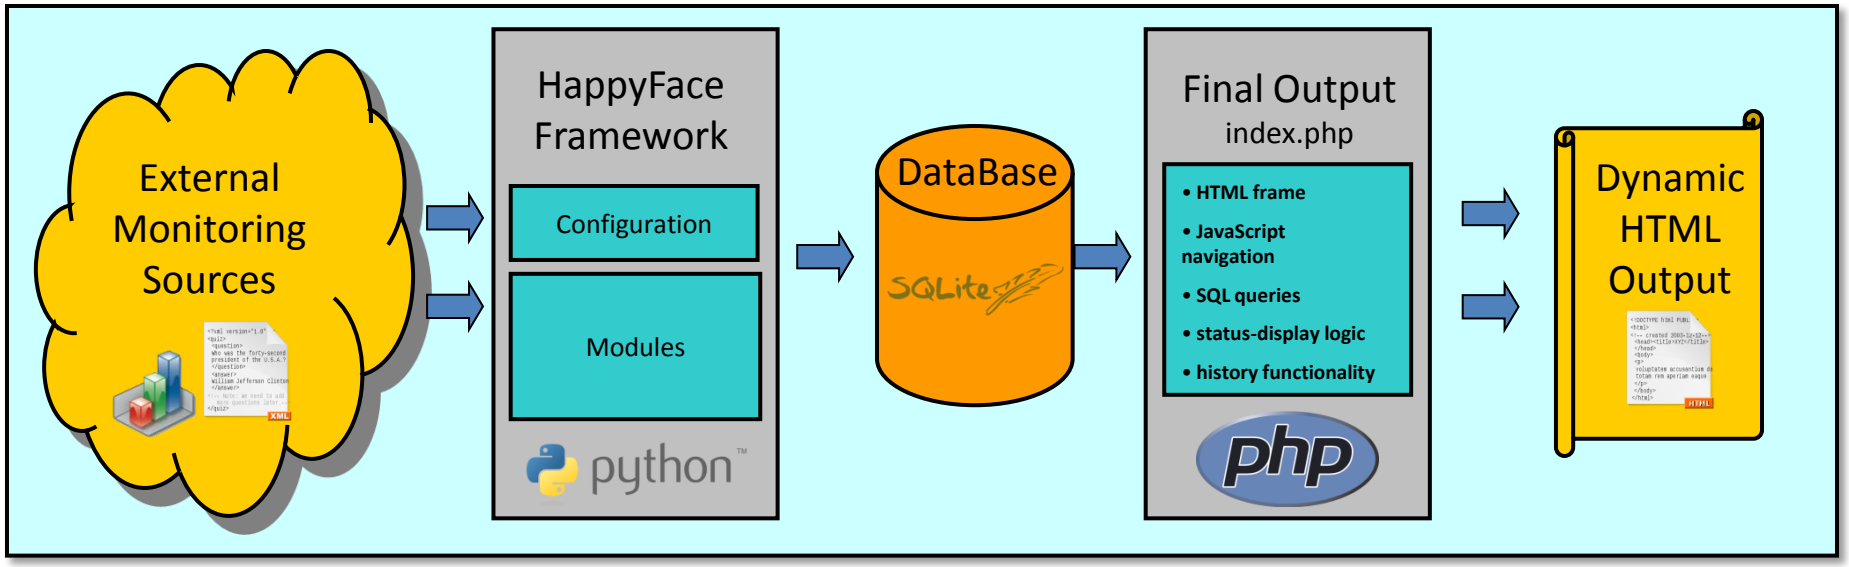

- The HappyFace Core provides all basic functionality needed by all tests and organises the test execution
- Each test is represented by a module, which can be plugged in
- Each module can be activated/arranged in the global configuration
- Core and all modules available on a central subversion repository

Each module …

- Can specify files to be downloaded
- Processes the gathered information
- Stores the output in the DataBase
- Provides a php fragment for the final web page

Configuration:

- Default configuration (which works out of the box) as .cfg file (in the SVN)
- Can be adapted for local purposes by a .local file

The HappyFace Webpage

Site logo

SVN revision

Stepwise history

Jump to certain point in time

Display most recent information

Categories: Name & Status

Category navigation bar

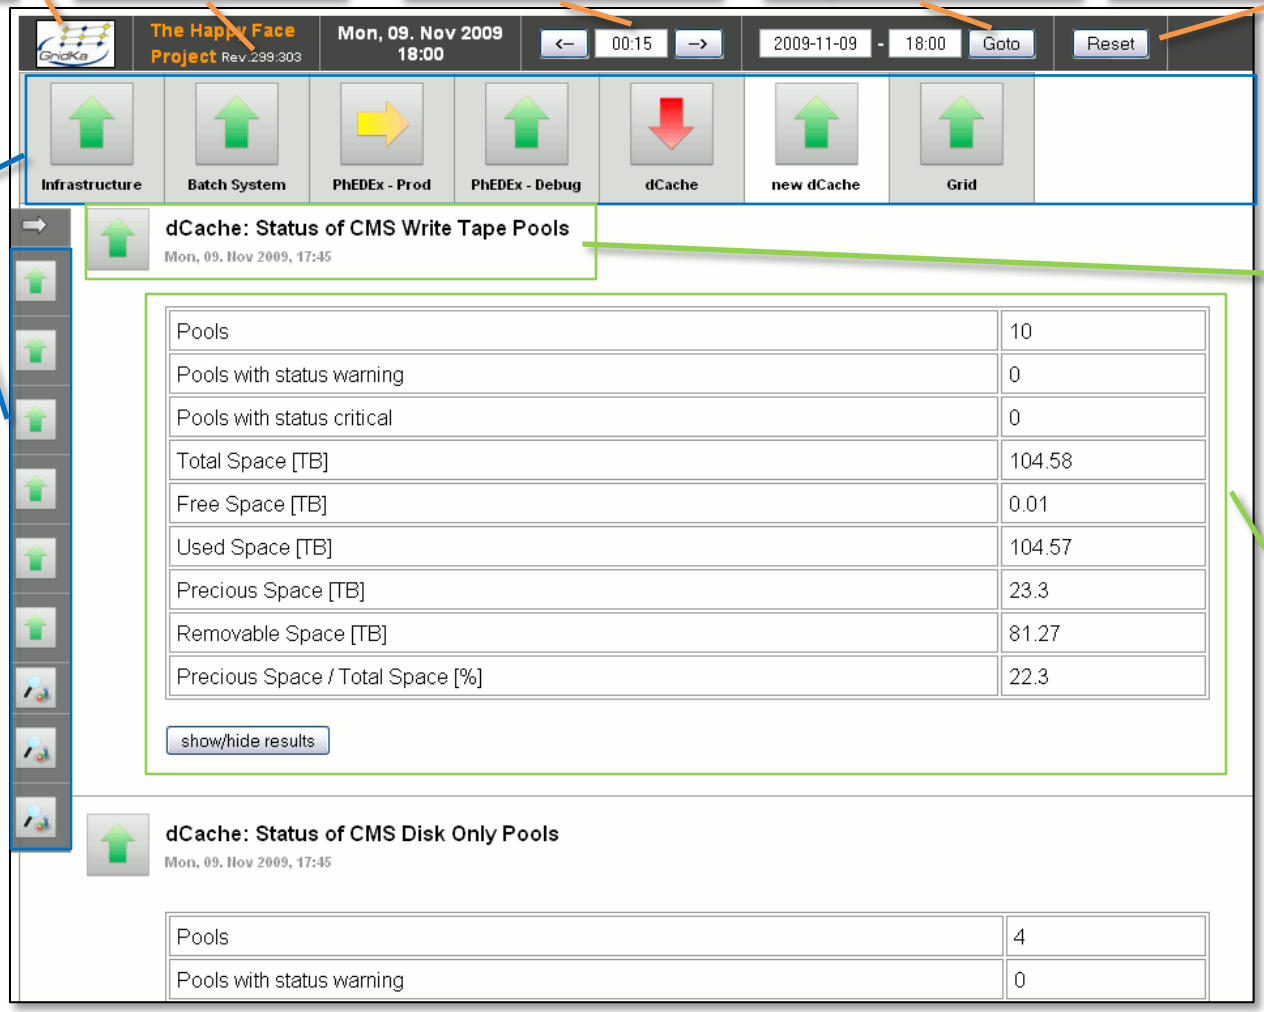

Module: Status, Title and execution time

Field for module output (php fragment created by the module)

A Closer Look at a Module

Click on the arrow to open the grey info box below

Unique name of the module

Is it rated or simply a plot?

Status of the module:
All ok (1.0) or error (0.0)

Weight for the category status calculation

Text fields for detailed module information

Create history plot for a given variable and time range

Open a sub frame with additional information (here status information of each pool)

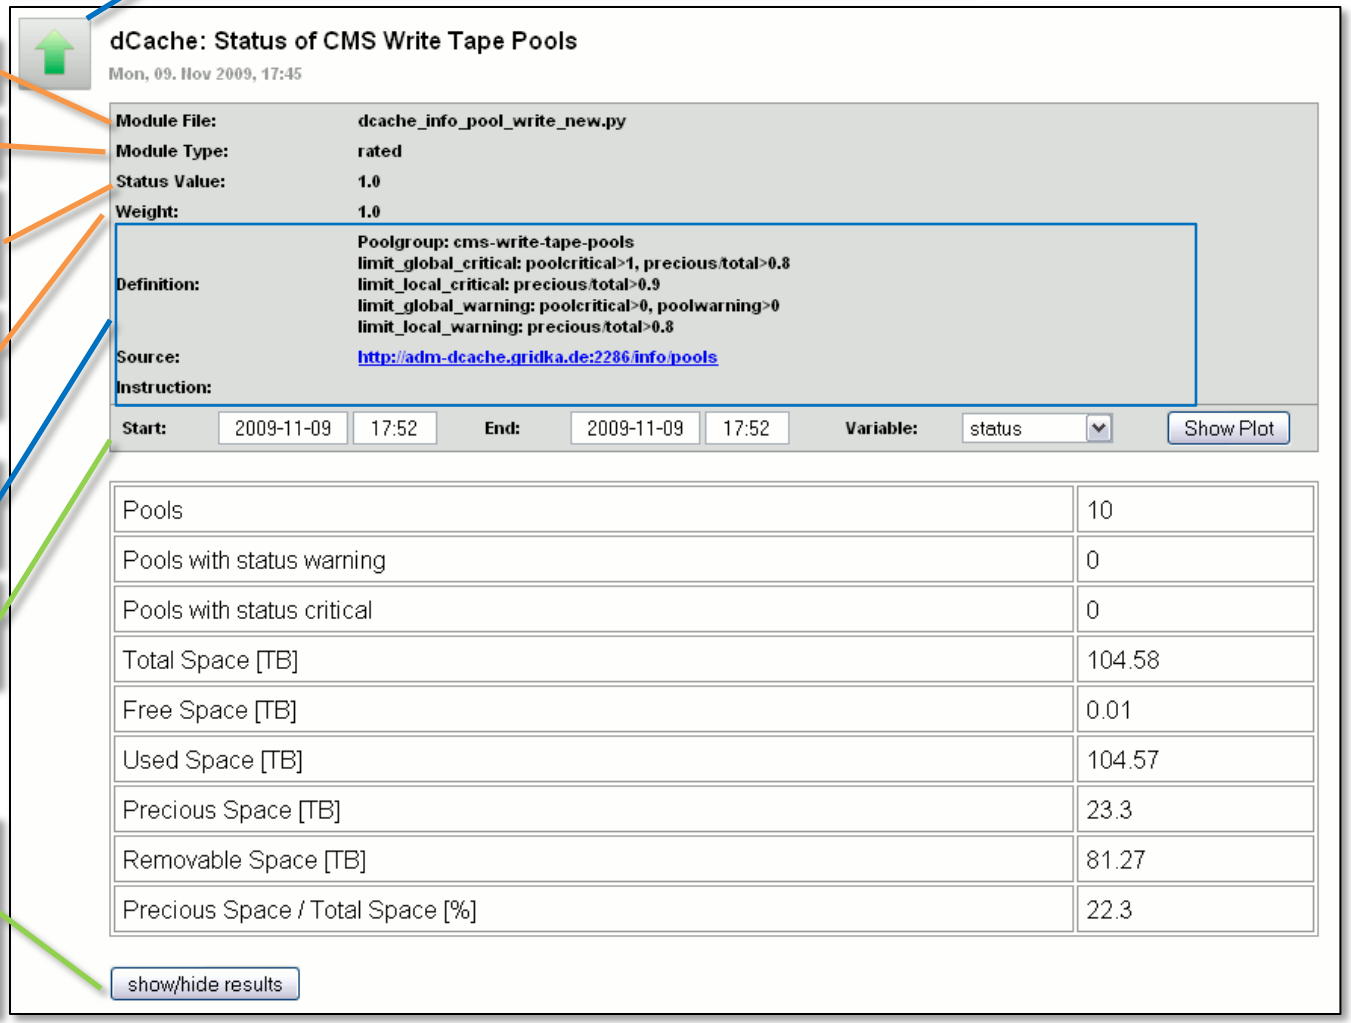

History

Module: dcache_datarestore_lazy_old

Variable: total

Start: 2009-11-09 18:03 End: 2009-11-10 18:03 Show Plot

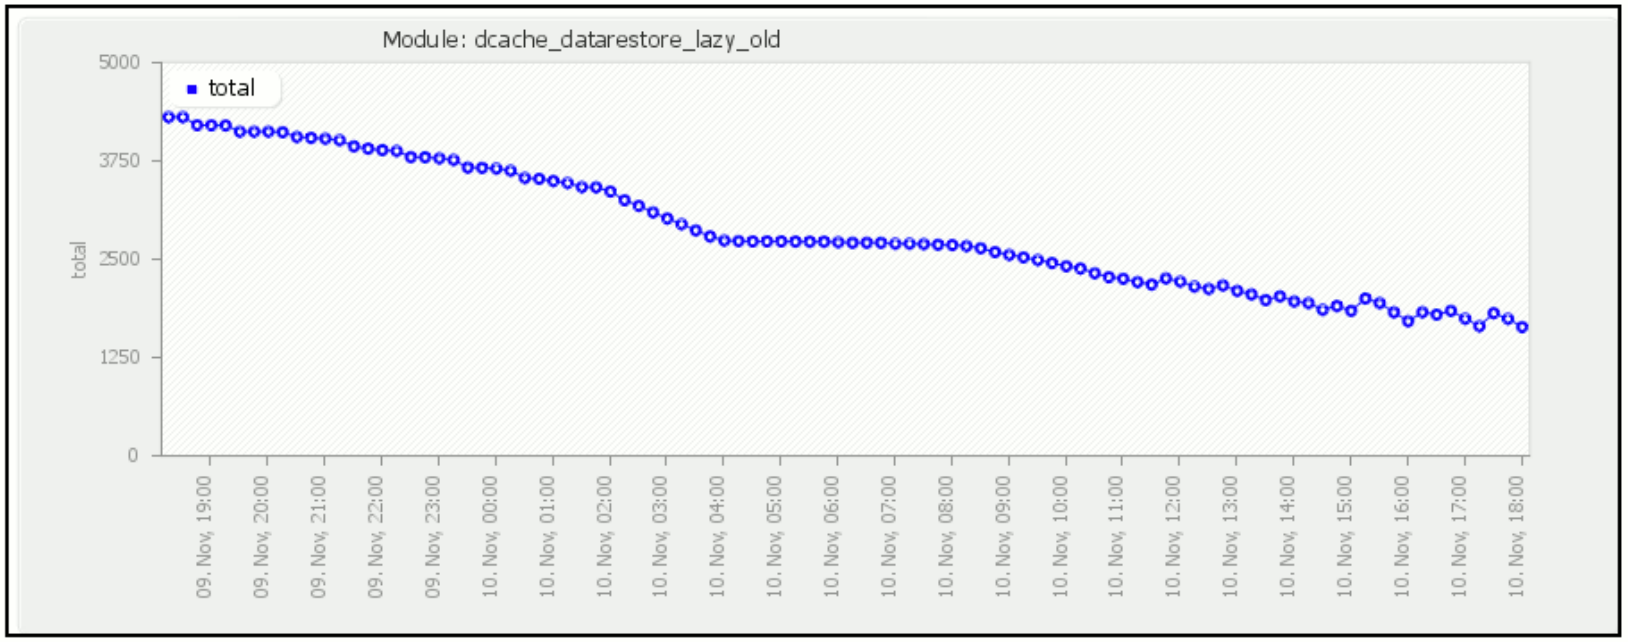

Module Architecture

- Each module is represented by a python class.
- All basic functionalities like database access, php output frame, history and other functions, which are needed by every module are provided by a basic class from which all modules inherit.
- Besides the initialisation, each module has two functions:
  - run(): collect the information, fill the database and determine status here
  - output(): define readout of the database and present the information
- Configuration:
  - All parameters of a module can be defined in a .cfg file. A default configuration is available in the SVN.
  - To configure a module for local particularities, the configuration parameters can be adapted via a .local file.
- Inheritance:
  - If more than one module uses a configuration parameter or a dedicated functionality: creation of a new class from which the modules inherit
  - Advantage: A whole bunch of modules can be adapted for a new site by only changing a parameter of one class (Example: URL and Site name for SAM information)

HappyFace Code and Instances

- Homepage of the HappyFace Project:
  https://ekptrac.physik.uni-karlsruhe.de/trac/HappyFace/wiki/Version_2

- System requirements:
  - Only a web server supporting PHP and Python is required! (and subversion for the code ;-))

- Development and maintenance:
  - HappyFace Core
    - **Karlsruhe:** Volker Büge, Viktor Mauch, Natalia Ratnikova and Armin Scheurer
  - The different modules:
    - **Aachen:** Philip Sauerland, Oleg Tsigenov

    https://lemon.physik.rwth-aachen.de/happyface
    - **Goettingen:** Stefan Birkholz, Joerg Meyer

    http://happyface-goegrid.gwdg.de/
    - **Hamburg:** Friederike Nowak

    http://wwwiexp.desy.de/groups/cms/tier2_monitoring/HappyFaceV2/trunk/webpage/index.php
    - **Karlsruhe**

    http://www-ekp.physik.uni-karlsruhe.de/~happyface/gridka/webpage/

Modules Karlsruhe I

dCache Pool Group Information

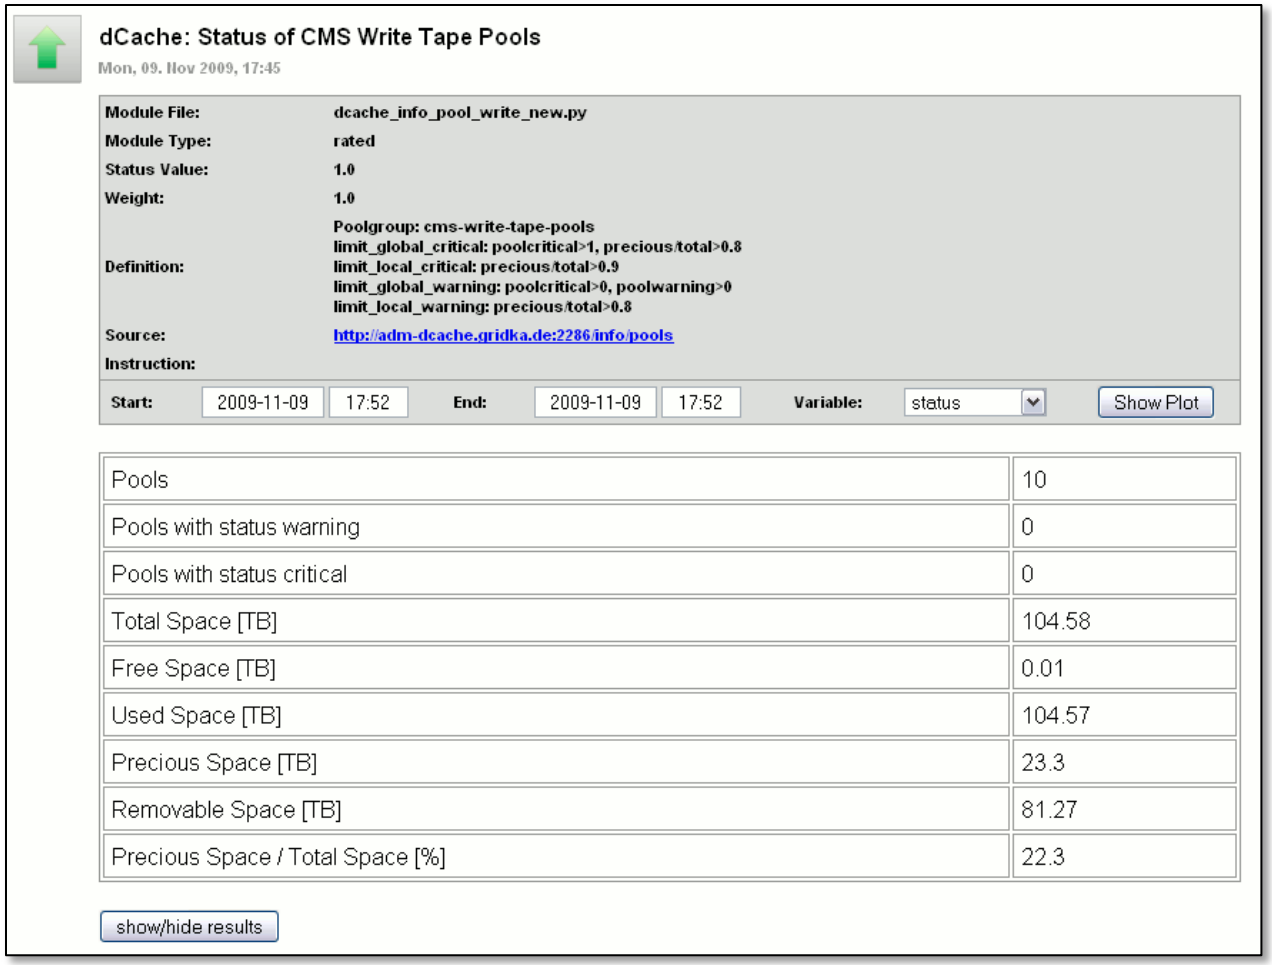

dCache: Status of CMS Write Tape Pools

Mon, 09. Nov 2009, 17:45

Module File:	dcache_info_pool_write_new.py
Module Type:	rated
Status Value:	1.0
Weight:	1.0
Definition:	Poolgroup: cms-write-tape-pools limit_global_critical: poolcritical>1, precious/total>0.8 limit_local_critical: precious/total>0.9 limit_global_warning: poolcritical>0, poolwarning>0 limit_local_warning: precious/total>0.8
Source:	http://adm-dcache.gridka.de:2286/info/pools
Instruction:	

Start: 2009-11-09 17:52 End: 2009-11-09 17:52 Variable: status Show Plot

Pools	10
Pools with status warning	0
Pools with status critical	0
Total Space [TB]	104.58
Free Space [TB]	0.01
Used Space [TB]	104.57
Precious Space [TB]	23.3
Removable Space [TB]	81.27
Precious Space / Total Space [%]	22.3

show/hide results

- Processing of the dCache Pool Information, provided by xml
- Select all information for a specified pool group (here cms-write-tape-pools)
- Summary of usage information, including the possibility to define thresholds for the module status
  - Ok, warning , critical
- Sub table with detailed information of each pool

Modules Karlsruhe II

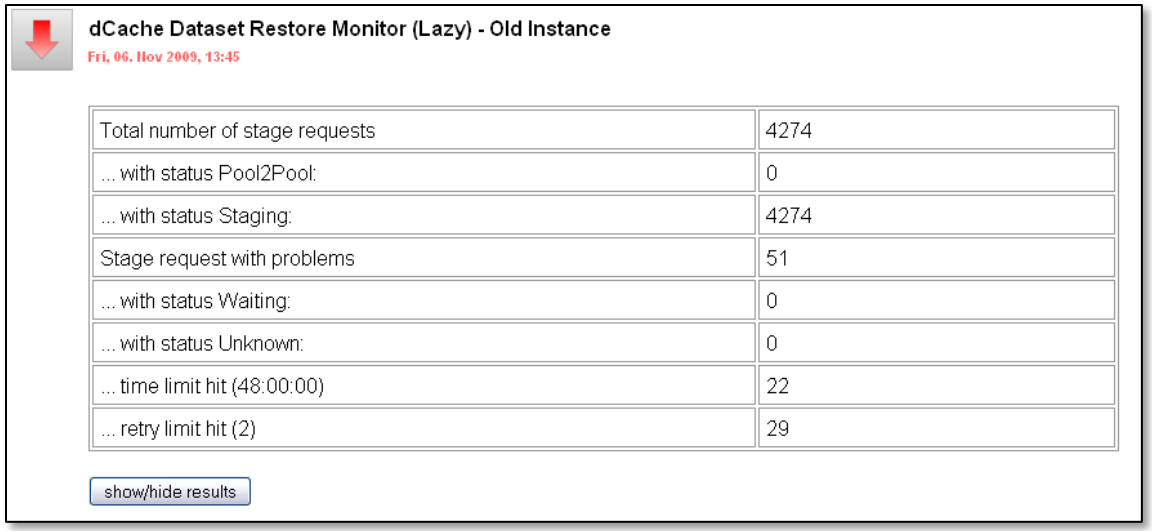

dCache Dataset Restore Monitor (Lazy) - Old Instance

Fri, 06. Nov 2009, 13:45

Total number of stage requests	4274
... with status Pool2Pool:	0
... with status Staging:	4274
Stage request with problems	51
... with status Waiting:	0
... with status Unknown:	0
... time limit hit (48:00:00)	22
... retry limit hit (2)	29

show/hide results

dCache Dataset Restore (Lazy)

- Processing of the dCache Dataset Restore Monitor web page
- Possibility to define thresholds of staging requests with problems
  - Time limit hit
  - Retry limit hit
  - Status waiting

SAM Test Results

- Summary of the SAM tests for a site
- Supports experiment specific and ops test
- Sub tables for summary of test results

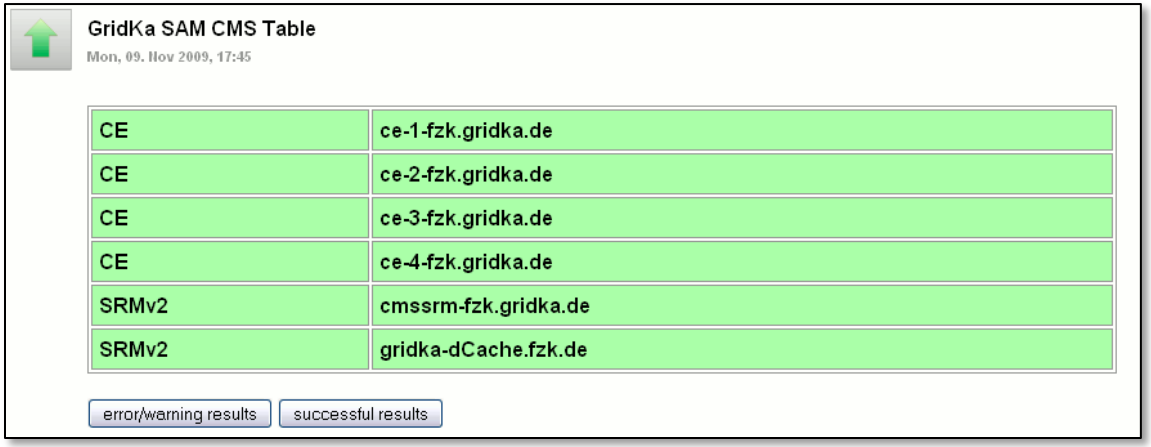

GridKa SAM CMS Table

Mon, 09. Nov 2009, 17:45

CE	ce-1-fzk.gridka.de
CE	ce-2-fzk.gridka.de
CE	ce-3-fzk.gridka.de
CE	ce-4-fzk.gridka.de
SRMv2	cmssrm-fzk.gridka.de
SRMv2	gridka-dCache.fzk.de

error/warning results | successful results

Modules Aachen

T2 User Space Monitoring

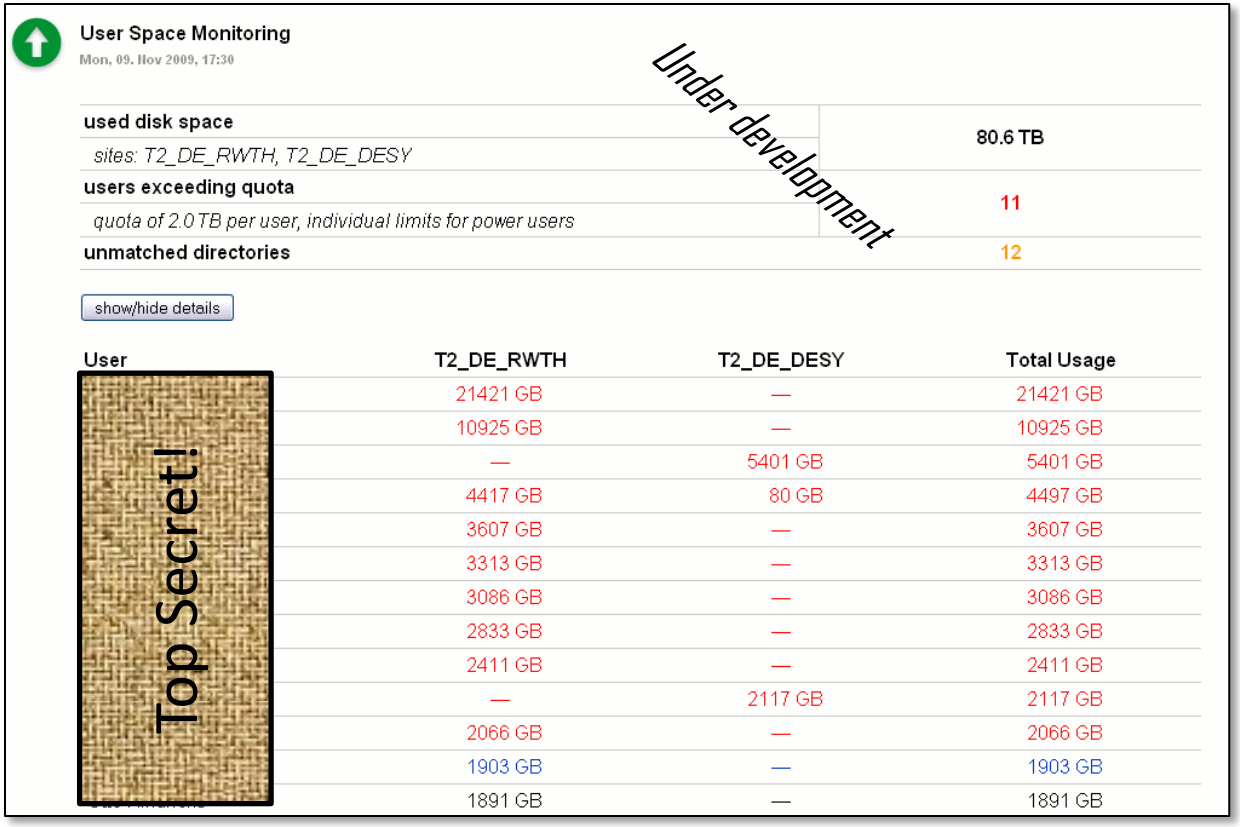

User Space Monitoring

Mon, 09. Nov 2009, 17:30

Under development

used disk space	80.6 TB
sites: T2_DE_RWTH, T2_DE_DESY	
users exceeding quota	11
quota of 2.0 TB per user, individual limits for power users	
unmatched directories	12

show/hide details

User	T2_DE_RWTH	T2_DE_DESY	Total Usage
Top Secret!	21421 GB	—	21421 GB
	10925 GB	—	10925 GB
	—	5401 GB	5401 GB
	4417 GB	80 GB	4497 GB
	3607 GB	—	3607 GB
	3313 GB	—	3313 GB
	3086 GB	—	3086 GB
	2833 GB	—	2833 GB
	2411 GB	—	2411 GB
	—	2117 GB	2117 GB
	2066 GB	—	2066 GB
	1903 GB	—	1903 GB
	1891 GB	—	1891 GB

- Information about used disk space per user exported via xml
- The HappyFace module reads in and processes these xml files per site
- Currently, database backend under development. The final version of this module will be released soon.
- Plan: Provide certificate based access to this information

Modules Hamburg I

CMS Phedex Agents

Phedex Agents (prod)

Mon, 09. Nov 2009, 17:15

agent	label	last report
BlockDownloadVerify	blockverify	00d:00h:04m
FileDownload	download-fts-all	00d:00h:27m
FileExport	exp-pfn	00d:00h:44m
FileIssue	mgmt-issue	00d:00h:01m
FilePump	mgmt-pump	00d:00h:03m
FileRemove	download-remove	00d:00h:00m
FileRouter	mgmt-router	00d:00h:01m

details

agent	label	host	directory	version
BlockDownloadVerify	blockverify	t2-cms-vo.desy.de	/home/cmssgm000/Prod_T2_DE_DESY/state	PHEDEX_3_2_9
FileDownload	download-fts-all	t2-cms-vo.desy.de	/home/cmssgm000/Prod_T2_DE_DESY/state	PHEDEX_3_2_9
FileExport	exp-pfn	t2-cms-vo.desy.de	/home/cmssgm000/Prod_T2_DE_DESY/state	PHEDEX_3_2_9
FileIssue	mgmt-issue	vocms01.cern.ch	/data/ProdNodes/Prod_Mgmt/state	PHEDEX_3_2_9
FilePump	mgmt-pump	vocms01.cern.ch	/data/ProdNodes/Prod_Mgmt/state	PHEDEX_3_2_9
FileRemove	download-remove	t2-cms-vo.desy.de	/home/cmssgm000/Prod_T2_DE_DESY/state	PHEDEX_3_2_9
FileRouter	mgmt-router	vocms01.cern.ch	/data/ProdNodes/Prod_Mgmt/state	PHEDEX_3_2_9

- Parses the XML provided by the PhEDEx server
- An agent not reporting back in a certain time leads to a warning/error status
- Error and warning thresholds are fully configurable
- Can be adapted for a new site by simply changing the site name in the .local file

Modules Hamburg II

Site Readiness for T2_DE_DESY

Mon, 09. Nov 2009, 17:00

Site Readiness Status	R	R	R	R	R	R	R	R	R	R
Maintenance	Up	Up	Up	Up	Up	Up	Up	Up	Up	Up
Job Robot	100%	n/a	n/a	100%	100%	100%	97%	100%	99%	99%
SAM Availability	100%	100%	100%	96%	100%	100%	88%	100%	100%	100%
T2::uplinkT1s	9	9	9	9	9	9	9	9	9	9
T2::downlinkT1s	8	8	8	8	8	8	8	8	8	8
date	30	31	01	02	03	04	05	06	07	08
month			Nov							

CMS Site Readiness

- HTML web site provides information about the “site readiness” for CMS
- Module parses the HTML and processes the information
- Allows to define warning thresholds

CMS PhEDEx Transfer Errors

- Parses the XML provided by the PhEDEx server
- Module distinguishes between source, destination, transfer and unknown error types
- Detailed information provided as sub table
- Error/warning thresholds are fully configurable

Transfer Errors to T2_DE_DESY (prod)

Mon, 09. Nov 2009, 17:00

failed transfers	8
failed transfers details	
failed transfers due to destination	0
failed transfers due to source	8
failed transfers due to transfer	0
failed transfers due to unknown reasons	0
fraction of destination errors	0%
fraction of source errors	100%
fraction of transfer errors	0%

details

Modules Goettingen I

Nagios Interface

- Summarizes warnings and error messages of Nagios monitoring
- Combines advantages of Nagios (lots of modules, including modules from EGEE) and HappyFace (lightweight, clear)
- Communication via ssh

Firefox Status Bar

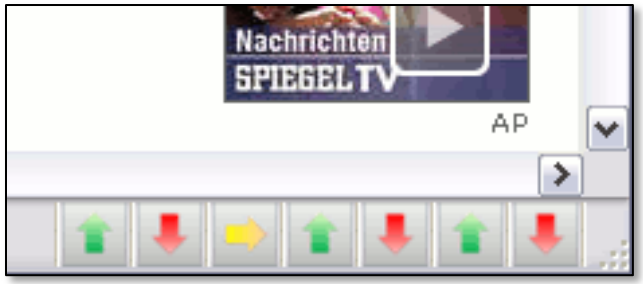

- HappyFace plug-in for Firefox browser available
- Summary of the category status of a site
- Clicking on an arrow gives the status information of the modules of a category

Modules Goettingen II

HP ILO Interface

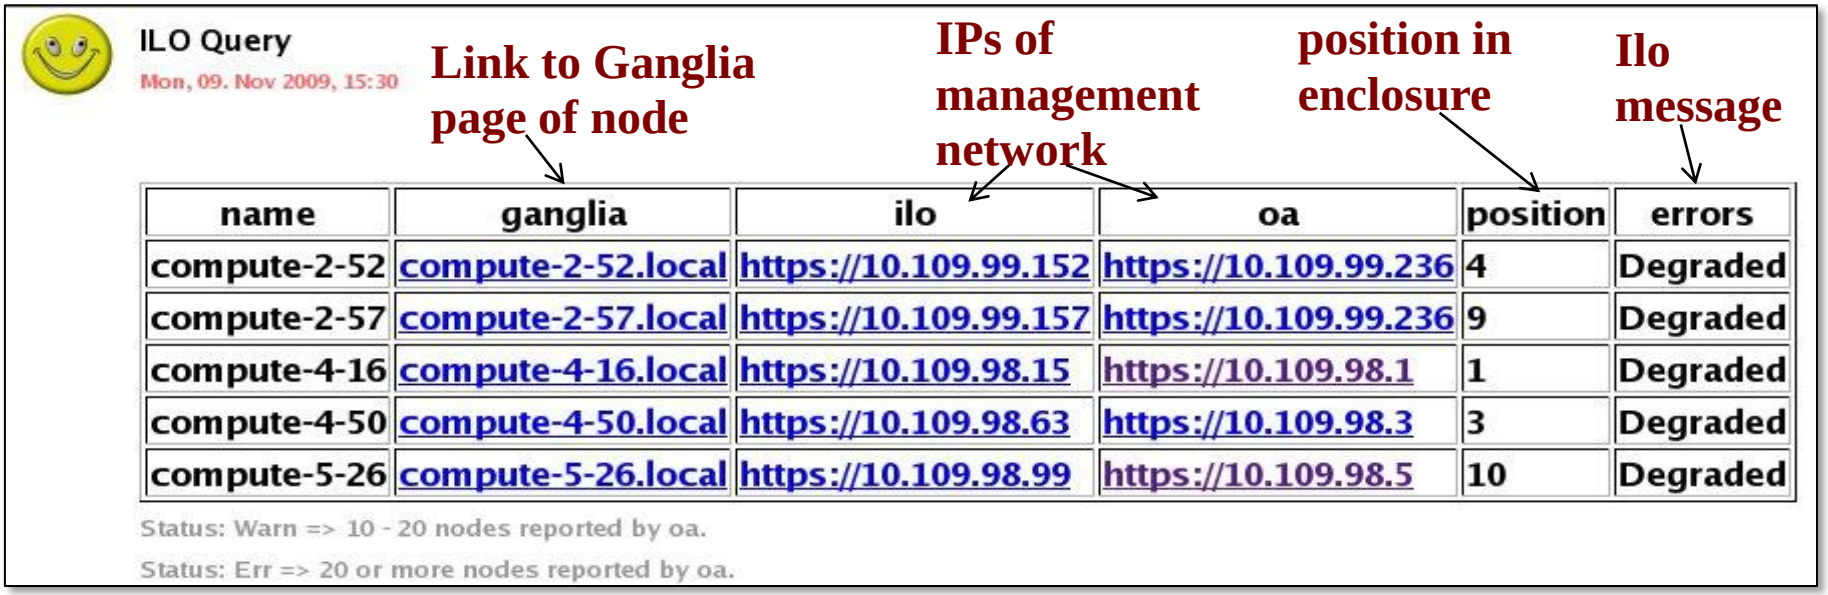

ILO Query

Mon, 09. Nov 2009, 15:30

Link to Ganglia page of node

IPs of management network

position in enclosure

Ilo message

name	ganglia	ilo	oa	position	errors
compute-2-52	compute-2-52.local	https://10.109.99.152	https://10.109.99.236	4	Degraded
compute-2-57	compute-2-57.local	https://10.109.99.157	https://10.109.99.236	9	Degraded
compute-4-16	compute-4-16.local	https://10.109.98.15	https://10.109.98.1	1	Degraded
compute-4-50	compute-4-50.local	https://10.109.98.63	https://10.109.98.3	3	Degraded
compute-5-26	compute-5-26.local	https://10.109.98.99	https://10.109.98.5	10	Degraded

Status: Warn => 10 - 20 nodes reported by oa.

Status: Err => 20 or more nodes reported by oa.

- Summarizes warnings and error messages of HP Integrated Lights Out (ILO) hardware monitoring
- Communication via ILO-XML-interface

… and this is only a selection of the available modules!

Conclusion & Outlook

The HappyFace framework:

- The HappyFace Project is a meta-monitoring suite, which provides a smart summary of existing monitoring sources
- It enables to link global with local monitoring information and to identify correlations
- It is not an additional monitoring information provider, it only collects and displays existing ones!

Organization:

- The core now exists in a stable release. A large variety of different modules is available for most use cases of WLCG T1 and T2 monitoring
- The development of new modules and features is taking place at all four HappyFace sites: Once a functionality is missing, it is simply added by the site which needs it!
- Number of partners is growing: The T1 GridKA will now use HappyFace for parts of their experiment specific monitoring

If you are interested in HappyFace, just have a look at the existing instances or download the package and give it a try.